

Intelligent, Self-Diagnostic Thermal Protection System for Future Spacecraft

Robert W. Hyers^{1,2}, Michael P. SanSoucie³, David Pepyne¹, Alaina B. Hanlon¹, and Abhijit Deshmukh¹.

¹University of Massachusetts, Amherst, MA 01003.

²corresponding author, hyers@ecs.umass.edu.

³ NASA Marshall Space Flight Center, Huntsville, AL 35812.

Abstract

The goal of this project is to provide self-diagnostic capabilities to the thermal protection systems (TPS) of future spacecraft. Self-diagnosis is especially important in thermal protection systems (TPS), where large numbers of parts must survive extreme conditions after weeks or years in space. In-service inspections of these systems are difficult or impossible, yet their reliability must be ensured before atmospheric entry. In fact, TPS represents the greatest risk factor after propulsion for any transatmospheric mission [1]. The concepts and much of the technology would be applicable not only to the Crew Exploration Vehicle (CEV), but also to ablative thermal protection for aerocapture and planetary exploration.

Monitoring a thermal protection system on a Shuttle-sized vehicle is a daunting task: there are more than 26,000 components whose integrity must be verified with very low rates of both missed faults and false positives. The large number of monitored components precludes conventional approaches based on centralized data collection over separate wires; a distributed approach is necessary to limit the power, mass, and volume of the health monitoring system. Distributed intelligence with self-diagnosis further improves capability, scalability, robustness, and reliability of the monitoring subsystem. A distributed system of intelligent sensors can provide an assurance of the integrity of the system, diagnosis of faults, and condition-based maintenance, all with provable bounds on errors.

1. Background

On the morning of February 1, 2003, the Space Shuttle Columbia broke up on re-entry. The Columbia Accident Investigation Board concluded that the cause of the accident was a piece of insulating foam that fell from the external tank 81.7 seconds after launch, striking the leading edge of the left wing and fracturing reinforced carbon-carbon (RCC) leading edge panel number eight [2]. The foam strike was not detected by the crew, nor observed from the ground until detailed review of video and photographs the next day [2]. Even after the foam strike was noticed, calculations suggested incorrectly that there was no cause for concern. Other methods of inspection were suggested, including taking photographs of the Shuttle from ground- and space-based cameras; however, detecting a small, black hole in a black panel against black space seems quite a challenge. During the 16 days the Shuttle was on orbit, "mission management failed to detect the weak signals that the Orbiter was in trouble and take corrective action" [2].

Detecting problems in the thermal protection system (TPS) is no small problem. The thermal protection system (TPS) is essential to protect the aluminum Shuttle from temperatures near 1640 °C experienced during reentry into the earth's atmosphere. The system is large and complex: the Space Shuttle orbiter is 37 m long with a 27 m wingspan and is covered by approximately 24,300 reusable tiles and 2,300 Flexible Insulation Blankets (FIB) [3]. The nose and the wing leading edges experience the highest temperatures upon reentry therefore Reinforced Carbon-Carbon (RCC) panels are used at these locations [3]. A total of 50 RCC panels are found on the Orbiter.

The thermal protection system represents the greatest risk factor after propulsion for any transatmospheric mission [1]. Any damage to the TPS leaves the Space Shuttle vulnerable and could result in the loss of human life as what happened in the Columbia accident. Currently no system exists to notify the astronauts or ground control if the thermal protection system has been damaged.

The Columbia Accident Investigation Board recommends "on-orbit inspection ... of the Thermal Protection System" "before return to flight" [2] The research described below consists of the preliminary steps leading toward implementing a biomimetic distributed sensor network in the thermal protection system of the Space Shuttle fleet and succeeding spacecraft. Such a network would provide continuous, real-time, high-resolution data on the "health" of this all-important system. Early awareness of the damage to Columbia might have allowed its tragic loss to be avoided, either through transatlantic abort of the launch (standard procedure in the event of engine failure) or in-space repair of the damaged parts.

Distributed sensor networks are common in biological systems. For example, the human skin contains more than 2 million nerve endings, each of which is specialized to detect touch, temperature, or pain. Such a system requires cheap, compact, lightweight sensors and a network architecture that minimizes the required communication bandwidth among sensors (nerve cells) and between the sensors and the brain. The system of sensors and their network must provide sufficiently low latency for responses to occur on an appropriate timescale, and must provide sufficient spatial resolution to allow assessment of the severity and locality of the stimulus.

This research proposes to develop a similar system of sensors to detect damage to the thermal protection system (TPS) of future manned spacecraft. Such a system would increase reliability of the spacecraft by providing continuous real-time feedback on the state of the TPS. This involves adding a simple microprocessor and several emitter-detector pairs to each of the tens of thousands of tiles or other logical units comprising the TPS covering the vehicle. Each tile then becomes a network node, with semi-autonomous capabilities.

2. Smart Tile Concept

In the current concept, the baseline capabilities of each tile consist of: detection of fracture within itself, optical temperature measurement at different depths within the tile, communication with neighboring tiles, and detection of loss of communication. These

capabilities may be implemented using fiber optics or other technologies. "Smart Tile" technologies are readily adaptable to both tile-based and monolithic ablative TPS. Adjacent regions may be considered "logical tiles", each with independent sensors, controllers, and communications. For the rest of the paper, we will use the term "tiles" to include both physical tiles (in a tiled TPS) and "logical tiles" in a TPS with continuous components.

The added hardware to each tile consists of a processor, sensors, communications, and power. The processors will be commercial off-the-shelf microcontrollers, which may be modified for radiation hardening if a suitable rad-hard part cannot be identified. With the quantities required, significant economies of scale will be available. These parts are very low power (hundreds of microwatts or less), so the entire system of tens of thousands of tiles will require only a few tens of watts. Since the communications occur over a distance of perhaps 20 centimeters, the optical communication may also be very simple and low-power. The weight added to each tile is less than a gram, so the total TPS weight increases by a few kilograms.

Different sensors are appropriate for different TPS materials and applications. Compatibility issues include service temperature, thermal expansion, and chemical reactivity. Temperature and integrity/fracture are likely to be important in all TPS applications, but some sensors such as recession are only relevant to ablative TPS. Once smart tiles are in place, their processors and communications network can support addition of a wide variety of potential sensors to any or all of the tiles. Such additional sensors may include active and/or passive acoustic sensors for monitoring components that are not themselves "smart".

The simplest sensors consist of an emitter-detector pair coupled to a continuous optical fiber. Continuity of this sensor fiber is lost if the tile is broken. Temperature monitoring by optical pyrometry may be achieved through the same fiber, allowing post-flight analysis of the TPS performance for more effective and efficient maintenance of reusable vehicles. The processor, emitters, and detectors will be permanently attached to the inside surface of the tile, next to the hull where the temperature remains moderate, $<150^{\circ}\text{C}$.

Adding intelligence to the thermal protection system involves many challenges that must be addressed and overcome in order to make the project a success. One important consideration is limiting the mass of the added hardware. Figure 1 shows the components that must be added to each tile. The orange sensors have the capability of monitoring the temperature at specific tile depths. The grey communication links found along the bottom of the tile serve as tile-to-tile communication lines that report to the processor found at the center of the tile's base. The green sensor is a continuous optical fiber that runs throughout the tile as shown, is coupled to an emitter-detector, and tests for tile integrity. In actuality, the fiber will have gradual curves rather than the sharp bends pictured in Figure 1.




Figure 1: "Intelligent" tile containing sensors and controllers: tile integrity (green), temperature (orange), processor (black), and communications (gray).

Failure of the RCC leading edge panel number eight was the proximate cause of loss of Columbia [2]. The health of RCC panels must be monitored in the future; however, adding "intelligence" to the RCC panels is a special challenge. Owing to their high service temperatures and high thermal conductivity, RCC panels do not provide conditions where current electronics may be safely housed, even on the inner surface of the panels. The solution which currently seems most promising will be investigated: using fiber detectors which pass through the RCC panels but are connected to emitter/detector pairs in the adjacent tiles. With this strategy, the tile-based network can also monitor the health of the RCC panels, while only the optical fibers are exposed to the extreme environment of the RCC leading edges and nose cone. Further investigation on the compatibility of optical fibers of quartz and other materials, with and without surface modification, with RCC materials at temperatures up to 1700 °C is necessary.

Similar capabilities are being considered for future ultra-high temperature ceramic (UHTC) leading edges. Such higher-temperature components may have only sensors installed, with processors and network communications removed to a more hospitable location nearby.

3. Smart Tile Sensors

Laboratory experiments were undertaken to demonstrate the feasibility of integrating fiber optics sensors into thermal protection system components. For this preliminary work, the Space Shuttle's silica TPS tiles were chosen for a number of reasons. The Shuttle tiles have well-characterized fabrication methods and well-characterized performance, allowing for comparison between laboratory tiles and actual service data. The fabrication procedure for the Shuttle TPS tiles is publicly available [8-11], again, allowing for representative laboratory experiments. Finally, the TPS materials to be used in future spacecraft are not yet well defined: even the balance between ablative and reusable TPS components is not yet determined.

The following sections describe the procedure used to fabricate Space Shuttle TPS tiles for service, followed by the procedures used in the laboratory experiments, the testing performed, and the results of the experiments on silica TPS tiles with co-fired fiber-optic sensors.

3.1 Background on Space Shuttle Thermal Protection System

The primary goal of the Thermal Protection System (TPS) is to limit the peak space-to-earth entry temperature and heat loads [3]. The highest temperatures occur on the nose cone and the leading edges of the wings. These areas are protected with reinforced carbon-carbon composites (RCC). The surface temperatures in these regions may reach 1500 – 1650 °C [4, 5]. The largest portion of the Orbiter's TPS originally consisted of almost 31,000 ceramic tiles bonded to the shuttle using nylon felt strain insulator pads [4]. However, over time many of these tiles have been replaced by approximately 2,300 Flexible Insulation Blankets [6].

Each ceramic tile are made from either High Temperature Reusable Surface Insulation (HRSI) or Low Temperature Reusable Surface Insulation (LRSI). The High Temperature Reusable Surface Insulation (HRSI) consists of LI-900, LI-2200, or FRCI-12 material coated with black borosilicate glass for high emittance. HRSI is used in areas that receive temperatures between 650°C and 1260°C, which are typically the lower surfaces of the Shuttle. The LI-900 and LI-2200 materials are entirely made of silica with average densities of 9 lb/ft³ (140.16 kg/m³) and 22 lb/ft³ (352 kg/m³), respectively. Areas requiring greater strength, such as door hinges, use LI-2200 tiles [7]. FRCI-12 is a composite insulation composed of silica fibers and aluminum-borosilicate fibers with a density of 12 cubic pounds per cubic foot [8]. Low Temperature Reusable Surface Insulation (LRSI) is used in areas that receive temperatures between 400°C and 650°C [5], and it is made of LI-900 with a white borosilicate coating [7]. They are coated white to limit solar heating and are typically used on the upper surfaces. About 80% of all the tiles are LI-900 [4].

The silica tiles are made from a very high purity amorphous silica fiber approximately 1.2 to 4 microns in diameter and 1/8 inch long [4]. The fiber is manufactured by Johns Manville and called Q-Fiber [9]. The silica tiles are 93% void, which makes them excellent insulators, with thermal conductivities as low as 0.017-0.052 W/m*K. The amorphous silica gives the tiles a low coefficient of expansion as well as a low modulus, and this eliminates any problems with thermal-stress and thermal-shock [5]. Crystalline forms of silica have a coefficient of thermal expansion over thirty times higher than that of amorphous silica [10].

This requires only the purest materials to be used and great care must be taken to avoid contamination during processing. Any impurities or contaminations during the production can cause crystallization of the material [9]. The most important feature of the silica fiber is its chemical composition. The fiber must not have an impurity content greater than 0.3% and the total alkali and alkaline earth content must be less than 0.06% [11].

The manufacturing process is lengthy and tightly controlled. First, the fibers are washed in dilute hydrochloric acid and then rinsed with deionized water. Next a binder is added to the fibers and the mixture is blended with water in a V-blender for 30-60 minutes, while the pH is held at a constant 9.0 using ammonium hydroxide. The binder is made by suspending fumed silica and starch in deionized water and ammonium

hydroxide [9], although the binder may, in fact, not be necessary to achieve the required properties [12]. This is done with controlled amounts and ratios to maintain the correct pH. The blended mixture is then poured into a mold and rapidly pressed at 69-138 kPa [9].

The resultant tile is then placed in an oven and the temperature is raised at 11°C per hour until 150°C. The total drying time is 18 hours. The tiles are then immediately placed in a furnace that is free of alkali or alkaline earth oxide impurities. The temperature is raised at 150°C per hour or less until 1200-1300°C is reached [9]. After the tile is fired, it is x-rayed to test for voids and to assure that it is still amorphous. If it passes inspection, it moves to the next stages of manufacture and is machined to its desired shape [13].

3.2 Laboratory Experiments

Sample Preparation

The samples were prepared for preliminary testing from the same precursor material as used in actual TPS tiles, Johns Manville Q-Fiber. The Q-Fiber was washed with water and then pressed into a rectangular shape using two ceramic discs. In the preliminary tests, no binder was used; however, the binder may not be necessary [12]. Quartz fibers of various diameters were placed through the middle of some samples, while some samples without fibers were also produced as control specimens. The fibers were used with the as-delivered surface finish. A heat treatment profile similar to that used by Lockheed was used to create the samples.

A total of 32 samples of various sizes and fiber loadings were produced and tested.

Testing Procedure

Three-point bending tests were performed on samples approximately 3 cm x 5 cm x 0.5 cm using an Instron mechanical tester running under computer control. The test setup is pictured in Figures 2 and 3. Clearly visible are the sample, 3-point bend apparatus, and one emitter/detector pair for each of the two fibers in the sample. This sample contains the largest diameter of fibers tested, and was chosen for display because the fibers are visible in the photograph.

A load cell monitored the force on the crosshead, while a linear voltage displacement transducer (LVDT) monitored the position of the crosshead. These outputs were connected to a PC-based data logging and control system.

The continuity of the fiber in the sample was monitored optically during the test. Each fiber was connected to a light emitting diode (LED) on one end and to a silicon PIN photodiode on the other end. Beads of quartz glass were melted onto the ends of the fibers to improve the optical coupling to emitter and detector. The output of the photodiode was connected to an amplifier to produce a voltage proportional to the intensity of light transmitted through the fiber. The amplifier output voltage was

connected to an analog-to-digital converter board in the computer data acquisition system, and was recorded along with the force and displacement data.

The tests were performed under constant crosshead speed of 0.5 mm/sec until after fracture of the sample.




Figure 2: Three-point bend specimen and test fixture. The glass fibers carry light from the LED's (blue mount) to the photodiodes (black mount) while the fibers remain unbroken. The intensity of the transmitted light is monitored by the wires from the photodiode.




Figure 3: Test sample after fracture. The photodiode current indicates that the fibers break when the tile breaks.

Results and Discussion

The load and displacement of a typical test is plotted in Figure 4 with the photodiode output voltage of two fibers. The load displays a fairly sharp maximum as the bottom of the tile (loaded in tension) fractures, with a broad decay due to the fibrous nature of the matrix material. Comparison with the fiber continuity signal indicates that fiber 2 fractured just at the maximum load (as expected), while fiber 1 broke at a few millimeters more displacement. Many of the tests showed similar behavior. This result indicates the need for a better understanding of bonding between the fiber and matrix.

A series of tests shows no effect of even the largest detector fibers tested on the flexural strength of the composite specimens.

The fracture surface of one of the test specimens was imaged by scanning electron microscopy, and is pictured in Figure 5. The matrix of fine fibers of amorphous silica and a single large detector fiber are clearly visible. This detector fiber shows a very small pull-out length, approximately 0.2 mm, indicating very good bonding between the detector fiber and the matrix. However, the range of pull-out lengths observed in 32 tests is about 0.2 to 7 mm, so improvement in the strength and consistency of the fiber-matrix bond is essential.

32 tests on specimens of TPS tile material with and without detector fibers were performed. These tests showed that the detector fibers could accurately and continuously determine the structural integrity of the specimen. However, there was a large variation in the sample displacement at which the fibers broke, even within the same specimens. This variation demonstrates the need for further study of the interface between fiber and matrix.





Figure 4: Results of a typical three-point bend test conducted at a constant crosshead displacement rate of 0.5 mm/sec. The sample fractures at a force of 8.2 N and fiber 2 fractures at the same time. Fiber 1 fractures at a slightly larger displacement.



600μm 50X

Figure 5: Electron micrograph of the fracture surface of a sample tile after testing. This detector fiber was very well bonded to the matrix, although other samples show a wide variation.

4. Smart Tile Networks: Communication and Distributed Decision-making

The smart tile concept consists of a network of sensors and processors interconnected by communication links. The core tasks of the smart tile sensor network are the detection, verification, classification, and timely notification of tile damage, tile misalignment, and missing tiles. To detect tile damage each tile is instrumented with a processor and one or more of the optical temperature and tile fracture sensors described previously. Tile misalignment and missing tiles are detected via the optical communication links that interconnect the tile's processors. Besides transferring message packets, the communication links are designed to be able to detect and distinguish misaligned and missing tiles from communication failures. To reduce the incidence of false alarms, failures detected by the sensor network are first communicated locally to neighbor sensors, which coordinate to verify and classify the severity of the failure. A simple yet robust randomized message routing protocol is then used to route a damage summary report through a Network Access Point (NAP) to the vehicle's main computers, where the report is either stored or communicated to the vehicle's pilots and ground control, depending on its severity. To minimize missed detections the smart tile network performs continuous self-monitoring from power-up to power-down.

4.1 Smart Tile Network Architecture

The major components of the smart tile sensor network are schematically illustrated in Figure 6. Based on the current Space Shuttle design, Figure 6 shows a TPS consisting of two types of tiles (non-critical and critical) and RCC panels. As shown in Figure 6, non-critical tiles contain a processor and temperature/fracture sensors, critical tiles contain multiple redundant processors and temperature/fracture sensors, and RCC panels have temperature/fracture sensors only. Processors are connected to their immediate neighbors by communication links. We propose optical communication links, with optical fibers that are optically aligned, but not physically continuous across the tile boundaries. The innovation of this design is that by monitoring the performance of the communication system, it is possible to detect and distinguish misaligned and missing tiles from communication failures.

4.2 Communication and Routing

The communication links interconnect the processors to their immediate nearest neighbors to form a lattice topology. Interspersed throughout the lattice are also a number of network access points (NAPs) whose purpose is to collect the alerts being communicated by the smart tiles and monitor the health of the sensor network. The NAPs are where the vehicle's electronics infrastructure interfaces to the smart tile sensor network. The number of hull penetrations is therefore related to the number of NAPs.

Packet-Based Messaging.

We propose a simple packet based messaging protocol in which tile damage alerts and network health test probes are sent hop-by-hop between the processors and NAPs. Each message packet contains the message source address, destination address, message type, and message payload. Standard techniques (e.g., CRC) are used to detect and correct communication errors.




Figure 6: Schematic of the smart tile sensor network architecture.

Randomized Routing.

For communicating tile damage alerts from the tiles to a NAP, a simple randomized routing protocol can provide an efficient and robust method. In particular, by encoding each processor and NAP with its (x, y) coordinate location each processor then knows its relative position within the lattice. Then, for example, a processor at location (x, y) forwarding a message to location (x', y') , where $x' > x$ and $y' > y$, will forward the message to a randomly selected neighbor with a probability biased to move the message closer to its intended destination, i.e., to a neighbor with larger x or larger y location in the lattice.

With such a randomized routing protocol, messages are forwarded in the “right” direction most of the time, but the routing protocol also occasionally allows a message to be forwarded in the “wrong” direction, see Figure 7. Such a randomized routing scheme enhances reliability since it can ensure (1) that a message is not lost if it encounters damaged tiles or broken communication links that cannot communicate anymore, and (2) that a message eventually reaches its destination if a path exists to the destination. Moreover, this fault-tolerance is obtained without the need for routing tables or complicated routing table updates after processor or link failure.

The cost of occasionally forwarding a message in the “wrong” direction is a marginal increase in the expected number of hops needed to route a message from a given source location to a given destination location. For instance, in a 30×30 lattice (900 tiles), where 90% of the messages are forwarded in the “right” direction and 10% of the messages are forwarded in the “wrong” direction, the hop count is only 24% higher than the minimum number of hops (in an $N \times M$ lattice the minimal number of hops is $N+M-2$). This marginal reduction in performance is more than compensated by increased robustness. Considering again a 30×30 lattice and assuming a certain probability of

communication error, in which case the message has to be resent in a new direction, the expected hop count increases only slowly. In particular, with a 5% probability of communication error, the increase in the expected number of hops is only 5%. For 10, 15 and 20% probability of communication error the increase in hop count is 11, 18 and 25% respectively. Since the anticipated error probabilities for the communication network are well under 1%, the unrealistically high error probabilities in these calculations clearly illustrate the robustness of our randomized routing scheme.




Figure 7 Randomized message routing. Arrow length indicates probability, i.e., with small probability a message gets routed in the “wrong” direction.

4.3 Error detection (verification and classification)

The approach used for error detection is different for non-critical tiles and critical tiles/RCC panels. While damage or loss of a few non-critical tiles does not jeopardize the integrity of the TPS, damage or loss to just a single critical tile could greatly increase the risk of vehicle loss on reentry. Because of the difference in the importance, non-critical and critical tiles/panels are instrumented differently. Each non-critical tile is monitored by a single processor and a single temperature/fracture sensor. Each critical tile/panel, on the other hand, is monitored by multiple processors and multiple temperature/fracture sensors. Detected problems must be classified as maintenance issues or emergency issues. Given that a detected problem could result in mission abort, extremely high confidence in the sensor network is a necessity. Therefore, minimal false alarm probability and missed detection probability are the main objectives for any monitoring system. These objectives are achieved by sensor coordination and information fusion.

Information Fusion in Non-Critical Tiles.

There are two ways failures are detected in non-critical tiles. The first one is the self-detection of an internal error, meaning that a sensor detects damage in the tile or a failure in its communication or sensor capabilities. The second involves processors periodically checking their neighbors to see if they are “alive” and functioning properly. A detected failure in either case results in the election of an Alert Coordinator (AC). The role of the AC is to attempt to verify the failure by collecting further information. This is done by

instructing the other neighbors of the faulty tile to verify its status. The AC then fuses the information and sends a report to a NAP.

There are two reasons for using an AC instead of just sending every piece of information to a NAP. First, the information is locally verified and is sent as one information packet to a NAP. This helps to reduce network traffic and less evaluation work is needed after the information reaches the NAP. Second, the false alarm probability can be reduced significantly. For example, suppose an AC finds that it cannot communicate with one of its neighbors. The processor could wrongly conclude that the tile is missing. But by querying the tile's other neighbors the AC can reduce the probability that its communication failure is due to a missing tile. Specifically, if the probability of communication link failure $p = 0.01$, then without coordination the probability of falsely reporting a missing tile would equal $p = 0.01$, since the AC acting alone would not be able to distinguish a missing tile from a failed communication link. However, by merging the information from four neighbor tiles, the probability falsely concluding a tile is missing (assuming independent communication error) is equal to $p^4 = 1 \times 10^{-8}$; a four order of magnitude reduction in false alarm probability.

Information Fusion in Critical Tiles and Panels.

Detected problems with a critical tile or panel must be handled with very high priority as they may indicate a potential emergency. But also failures in communication links, sensors or processors have to be double checked, since they are crucial for the confidence of the report. To acquire additional information each critical tile and RCC panel has multiple redundant integrity sensors. When a failure is detected, one processor is elected as a Cluster Leader (CL). The CL may be a processor in the tile where the failure is detected, or as in the case of RCC panels, may be in an adjacent tile. A CL not only detects failures through its own sensor, but is also the point of contact for information gathered by the other sensors and processors in its cluster. Using, for example, simple majority vote, information on damage and failures is fused by the CL and reported to a NAP. The more independent sensors that are available, the more reliable the voting scheme is in reducing false alarms and missed detections. In case of a failure of the CL itself, a neighboring processor would detect the CL failure and would then take over the role as a CL.

Self-Monitoring.

The worst situation is to have an alarm system that has failed and to not know that the alarm is not functioning. Confidence in the smart tile sensor network, therefore, requires constant checking of its functionality. The central mechanism for doing this is for processors to periodically "ping" the status of their neighbors. Any detected errors are verified, classified, and communicated to a NAP as described previously. As an additional independent check of the lattice, the NAPs are not only recipients of messages from the processors, but the NAPs can also actively query individual tiles along selected routes through the lattice. In this way, the NAPs can obtain a check of the network's health and topology.

Additional Features

The speed and reliability of a routing mode can be further improved by sending messages to more than one NAP. This not only reduces latency, but also reduces the probability of message loss when large portions of the network are out of service due to tile damage or for some other reason (e.g., power failure). There is, of course, a trade-off between network traffic and number of NAPs to which the message is sent. Other extensions to the concept could include the recording of tile temperature histories for maintenance scheduling and possibly vehicle attitude control during reentry.

4.4 Computational Test Bed and Hardware-in-the-Loop Simulation

To test the proposed system a network of 26,000+ tiles with their processors, sensors, and communication links will be emulated using a grid-computing cluster. A variety of scenarios and topologies will be evaluated. In addition, the parameter settings for the randomized routing protocol will be optimized taking into consideration network bandwidth, message size, lattice topology, number and distribution of NAPs, expected number of tile failures and their expected spatial distribution as well as errors in message transmission. Finally, to test the hardware concept, we propose to manufacture ~10 tiles with real processors, sensors, and communication links and integrate them with the grid-computing cluster for hardware-in-the-loop test and evaluation.

5. Conclusions

In-situ assessment of the condition of the thermal protection system in near real-time is a challenging problem, but is an essential capability for future manned spacecraft. Smart materials and components such as those described here are a very promising part of the solution to this problem. Fiber-optic sensors meet many of the requirements for this application, including a high service temperature, low power, and low mass.

Distributed sensing with local intelligence and distributed decision-making is critical to meeting the constraints on mass, power, and volume for the entire condition-monitoring system. Many locations within and adjacent to the TPS have environments benign enough to support current electronics, and the range of suitable locations will increase with the availability of higher-temperature electronics.

In summary, smart tiles will play an important role in improving the safety, reliability, and serviceability of future spacecraft.

6. Acknowledgements

The authors wish to thank many people for their helpful and critical input on this work, including Chuck Smith, Raj Venkatapathy, Ed Martinez, and Frank Milos of NASA Ames Research Center, Space Technologies Division; Fanny Zuniga, Dougal MacLise, and Phil Wysocki of the NASA ARC Integrated Systems Health Management group; and Joy Huff, Lisa Huddleston, and Jennifer Gill of the Thermal Protection System Group at NASA Kennedy Space Center. We also wish to thank Al Rakouskas and Rich Bradshaw (UMass) for assistance with the experiments. This work was supported in part by a grant from the Massachusetts Space Grant Consortium.

7. References

1. Charles Smith, "Entry Systems Technology", Presented at NASA Ames Research Center Industry Day, July 22, 2004.
2. *Columbia Accident Investigation Board, Report Volume 1*, NASA, Washington, DC August 2003.
3. J.D. Buckley, G. Strouhal and J.J. Gangler, "Early Development of Ceramic Fiber Insulation for the Space Shuttle," *Ceramic Bulletin*, vol. 60, no. 11, 1981, pp. 1196-1200.
4. L. Korb and H. Clancy, "The Shuttle Orbiter Thermal Protection System: A Material and Structural Overview," *26th National SAMPE Symposium*, vol. April 28-30, 1980, pp.232-249.
5. L. Korb, C. Morant, R. Calland and C. Thatcher, "The Shuttle Orbiter Thermal Protection System," *Ceramic Bulletin*, vol. 60, no. 11, 1981, pp. 1188-1193.
6. http://science.ksc.nasa.gov/shuttle/technology/sts-newsref/sts_sys.html.
7. R.L. Dotts, D.M. Curry and D.J. Tillian, "Orbiter Thermal Protection System," *NASA CP 2342 Part 2*, 1983, pp. 1062-1081.
8. D. B. Leiser, M. Smith and H. E. Goldstein, "Developments in Fibrous Refractory Composite Insulation" *American Ceramic Society Bulletin*, vol. 60, no. 11, 1981, pp. 1201-1204.
9. H.E. Goldstein, M. Smith, and D. Leiser, "Reusable Silica Surface-Insulation Material," *NASA Tech Brief B73-10504*, December, 1973.
10. W. Schramm, "HRSI and LRSI - the Early Years," *Ceramic Bulletin*, vol. 60, no. 11, 1981, pp. 1194-1195.
11. J.C. Fletcher, H.E. Goldstein and D. Leiser, *Silica Reusable Surface Insulation* U.S. Patent 3,952,083, 1976.
12. D. Leiser, *Private Communication*, 2003.
13. J.M. Corp, *Private Communication*, 2003.